IMPLEMENTATION OF INTENSITY FRONTIER BEAM INFORMATION DATABASE

Igor Mandrichenko, FNAL

CHEP 2012, New York, NY

Mission

- The IFBeam Database system is responsible for extraction of data from the FNAL accelerator division systems and providing it to the different experiments involved in the Intensity Frontier programs.

Requirements

- Receive beams conditions data from the accelerator in real time
- Store information in the relational database
- Make the data available to online monitoring, data processing and analysis systems of the intensity frontier experiments
- Data loss must be minimized
- Acceptable data latency is 1 hour for long term data and ~minute for monitoring data
- Data preservation and recovery procedures
- Long term data should be stored forever
  - Used for data processing
- Short term data should be stored for several days
  - Used for online monitoring

Data Streams

- Slow data stream
  - All data (short term and long term) stored in the database
  - Buffered on the collector nodes and stored later by the Merger
  - Latency ~5 minutes
- Fast event stream
  - Only last event is stored is stored in the database
  - Stored by the Collector directly
  - Latency < 1 second
  - Can be lost
- Event timestamps
  - Sent as a UDP message from the Collector to the Web Server
  - Latency < 1 second
  - Can be lost

Database Schema

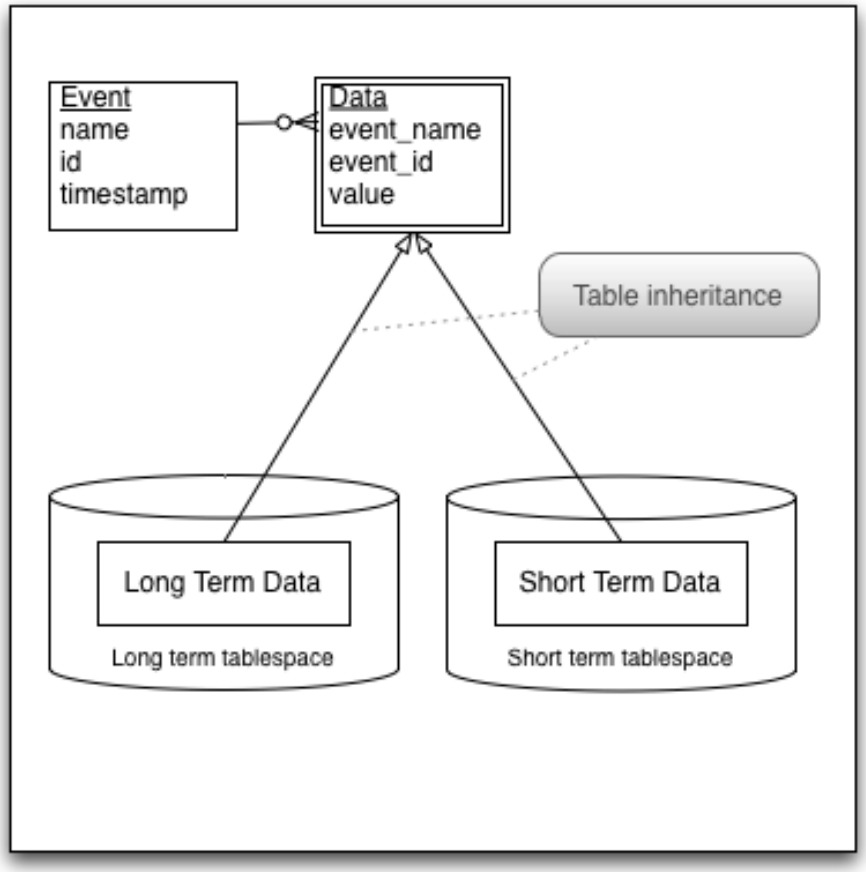

- Data tied to event via foreign key
- Short term and long term data stored in 2 different tables, inheriting from common “data” table
- Daily short term data cleanup uses “concrete” short term table
- Data access from “abstract” table

IFBeams Data Collection

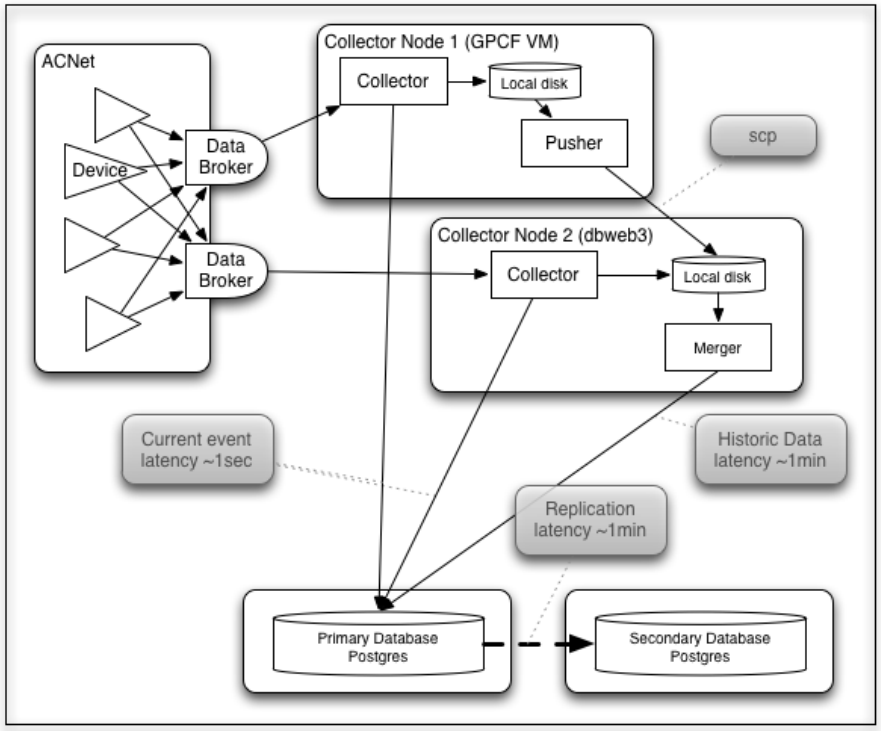

- Event rate – from 0.5Hz to 15Hz in peak
- Data per event
  - ~ 50 devices, ~1500 floating point numbers – long term
  - ~450 devices, ~2000 f.p. numbers – short term
- Current counts:
  - Long term data 400M
  - Short term data 20M
  - Events: 52M
- Current size:
  - 250 GB
    - mostly long term data (200GB, including 60GB index)
  - Estimated growth 0.5-1TB/year
  - Estimated end of life size ~10TB

Reliability of Data Collection

- Data is redundantly collected and buffered by 2 identical Collectors running on 2 different computers
  - Merger eliminates the redundancy and stores data into the database
  - More Collectors can be added
- Any one Collector can be turned off without any effect on the system
- Data can be buffered by Collectors for weeks in case of database or network outage

Data Access

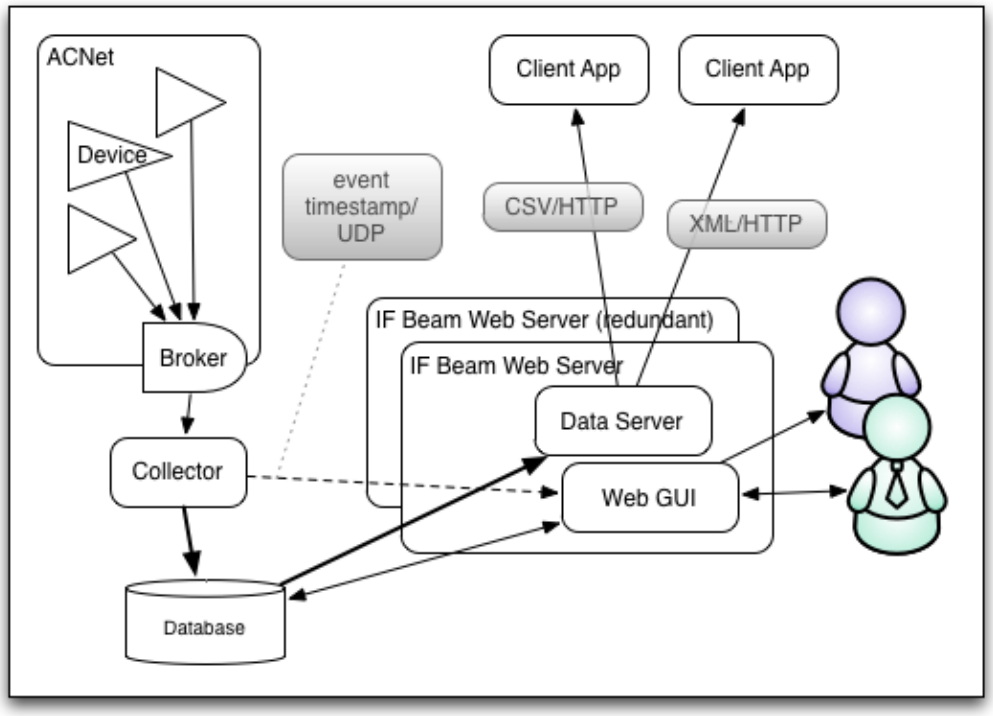

- Applications get data via REST/HTTP from the web data server
  - XML
  - CSV
  - JSON
- Redundant web server infrastructure increases availability
- Interactive Data Browser
- Dashboard
- Data collection controls

Reliability of Storage

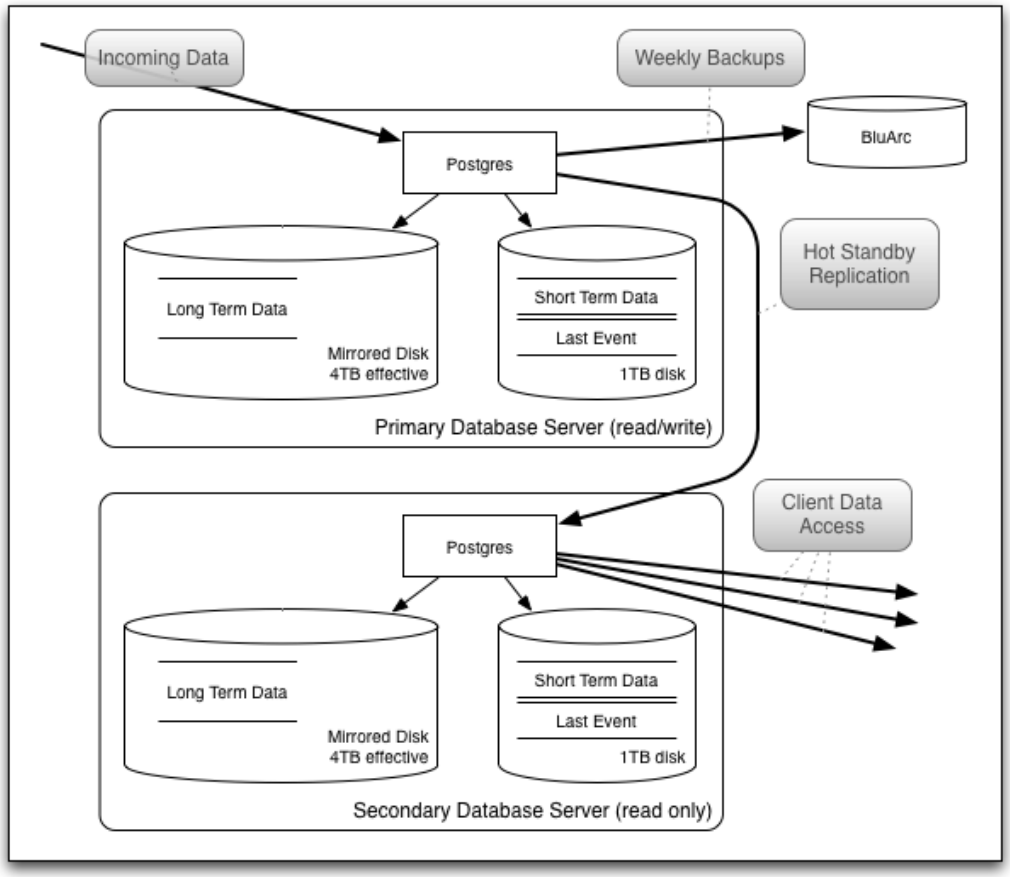

- Data is stored on 2 computers
- Postgres hot standby replication is used for replication
- Loss of any single computer is easily recoverable
- Disaster recovery: weekly backups to external network attached storage

Dashboard

- Used to monitor health of all the components of the system, from the collector to web data server
- Recorded event frequency
- Measurements from key devices